A Power-Efficient, Scalable Processor Family

Presented by Jim Keller

VP Engineering, Architecture Group

Fall Processor Forum

October 25th, 2005

PA SEMI

Power to Perform™

Outline

- Motivation
- Family
- Core
- SOC
- Performance and power
- Applications
- Summary

Motivation

- **65 nm opens up new possibilities**
  - Enables complete platform integration
  - But requires a new approach to power management
- **New generation of I/O converged on SERDES**
  - PCI Express, XAUI, SGMII, SATA, RapidIO…..
  - Opportunity to make a very versatile chip
- **Common goal**
  - More processing, I/O, functionality
  - Less power, parts, area, latency, overhead

Introducing PWRficient™

- **Multicore SOC family based on Power cores**
  - Architectural license from IBM
  - Ground-up implementation by P. A. Semi
- **Very broad application range**
  - Networking, storage, imaging, security applications
  - Game consoles through super servers
- **New design maximizes computing density and reliability**
  - Optimize performance vs. power vs. size
  - Integrate high-bandwidth I/O and memory
  - Focus on low latency, high throughput

PWRficient Family

PA SEMI

Power to Perform™

PWRficient Family

- **PA6T core—the CPU**
  - Power compliant, 32/64-bit, high performance FPU & VMX
  - 7W @ 2GHz worst-case power dissipation
  - 4W @ 1.5GHz worst-case power dissipation
- **CONEXIUM™—the on-chip coherent interconnect**
  - Scalable cross bar interconnect
  - 1–8 SMP cores
  - 1 or 2 L2 caches, sized 512KB–8MB
  - 1–4 64-bit DDR2 memory controllers
- **ENVOI™—the I/O system**
  - SERDES I/O—PCI Express, XAUI, SGMII, SATA/SAS, RapidIO
  - Offload engines—TCP/IP, iSCSI, cryptography, and RAID
  - Support I/O—Boot bus, UARTs, SMBus, GPIO

PWRficient 1682M Block Diagram

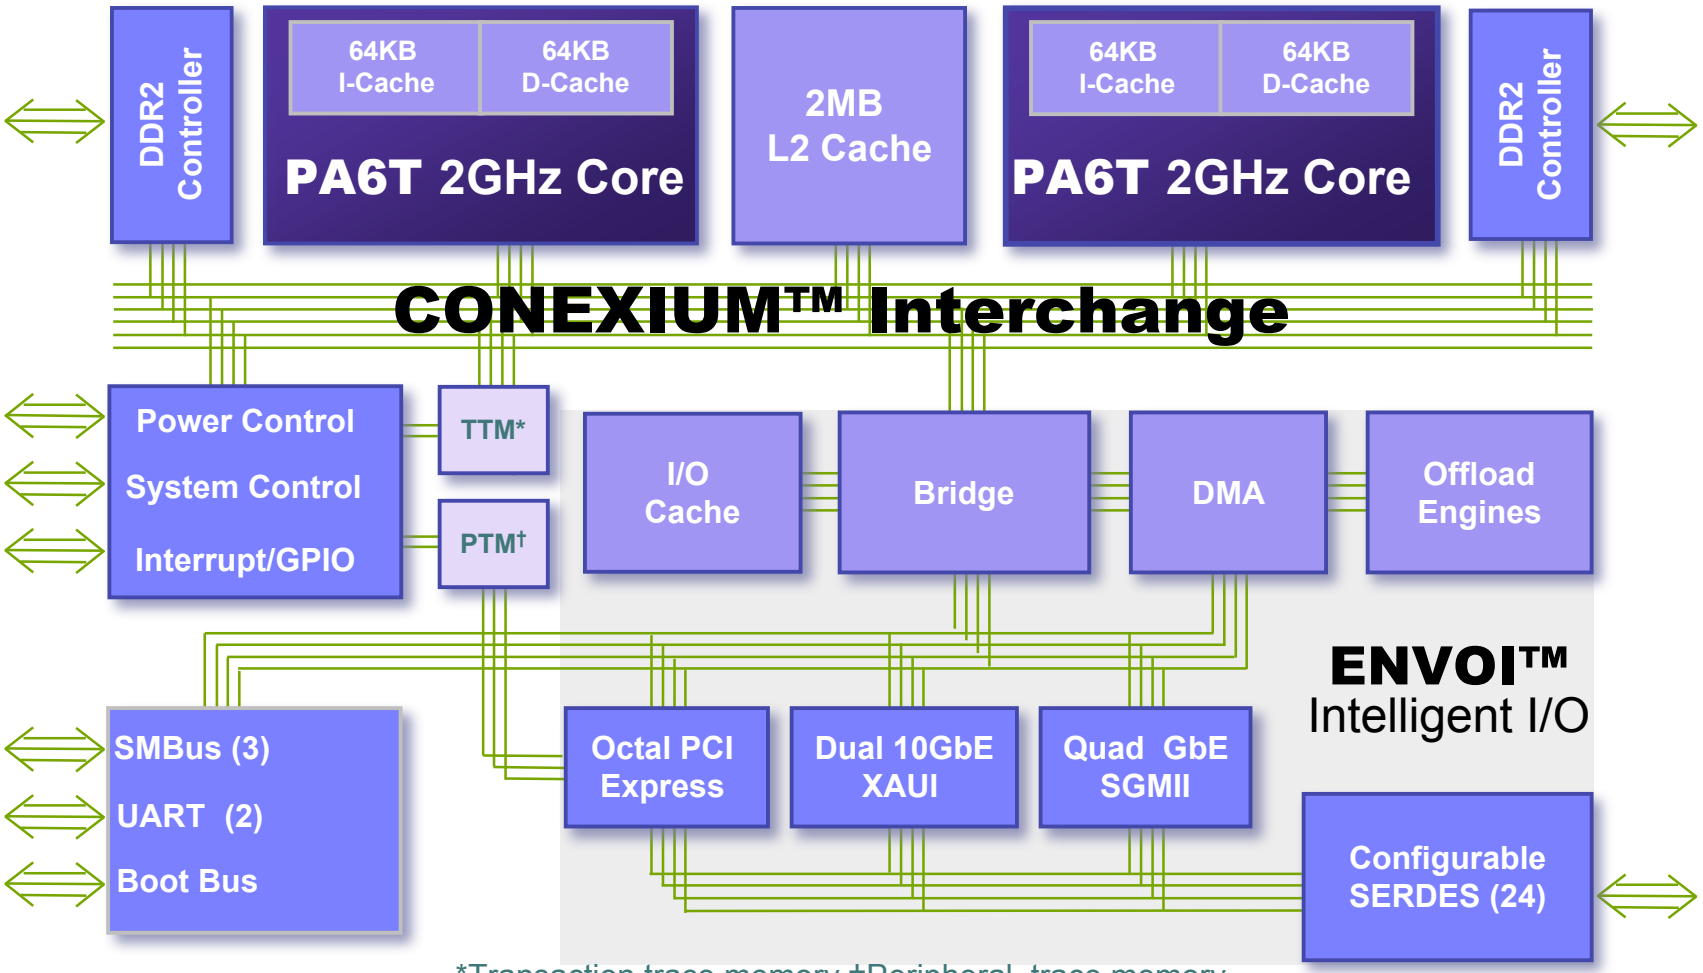

*Transaction trace memory †Peripheral trace memory

Scalable Sockets

Common sockets enable simple customer upgrade/cost reduction

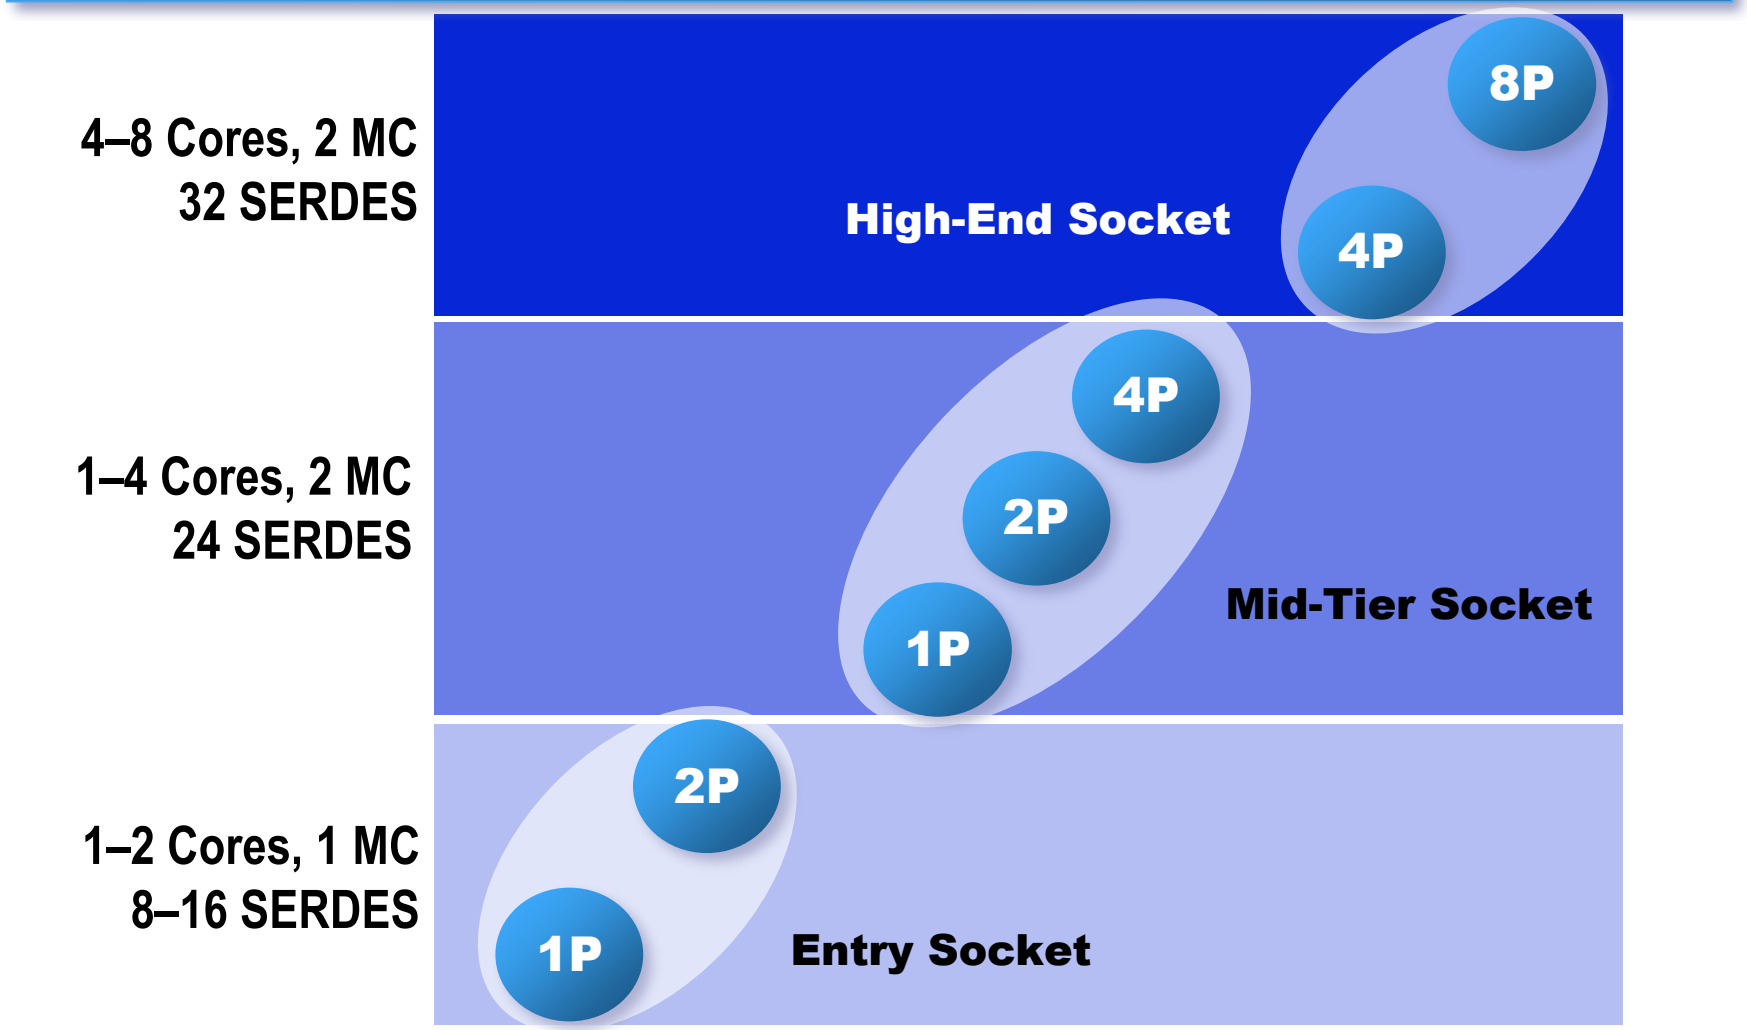

PWRficient Rollout*

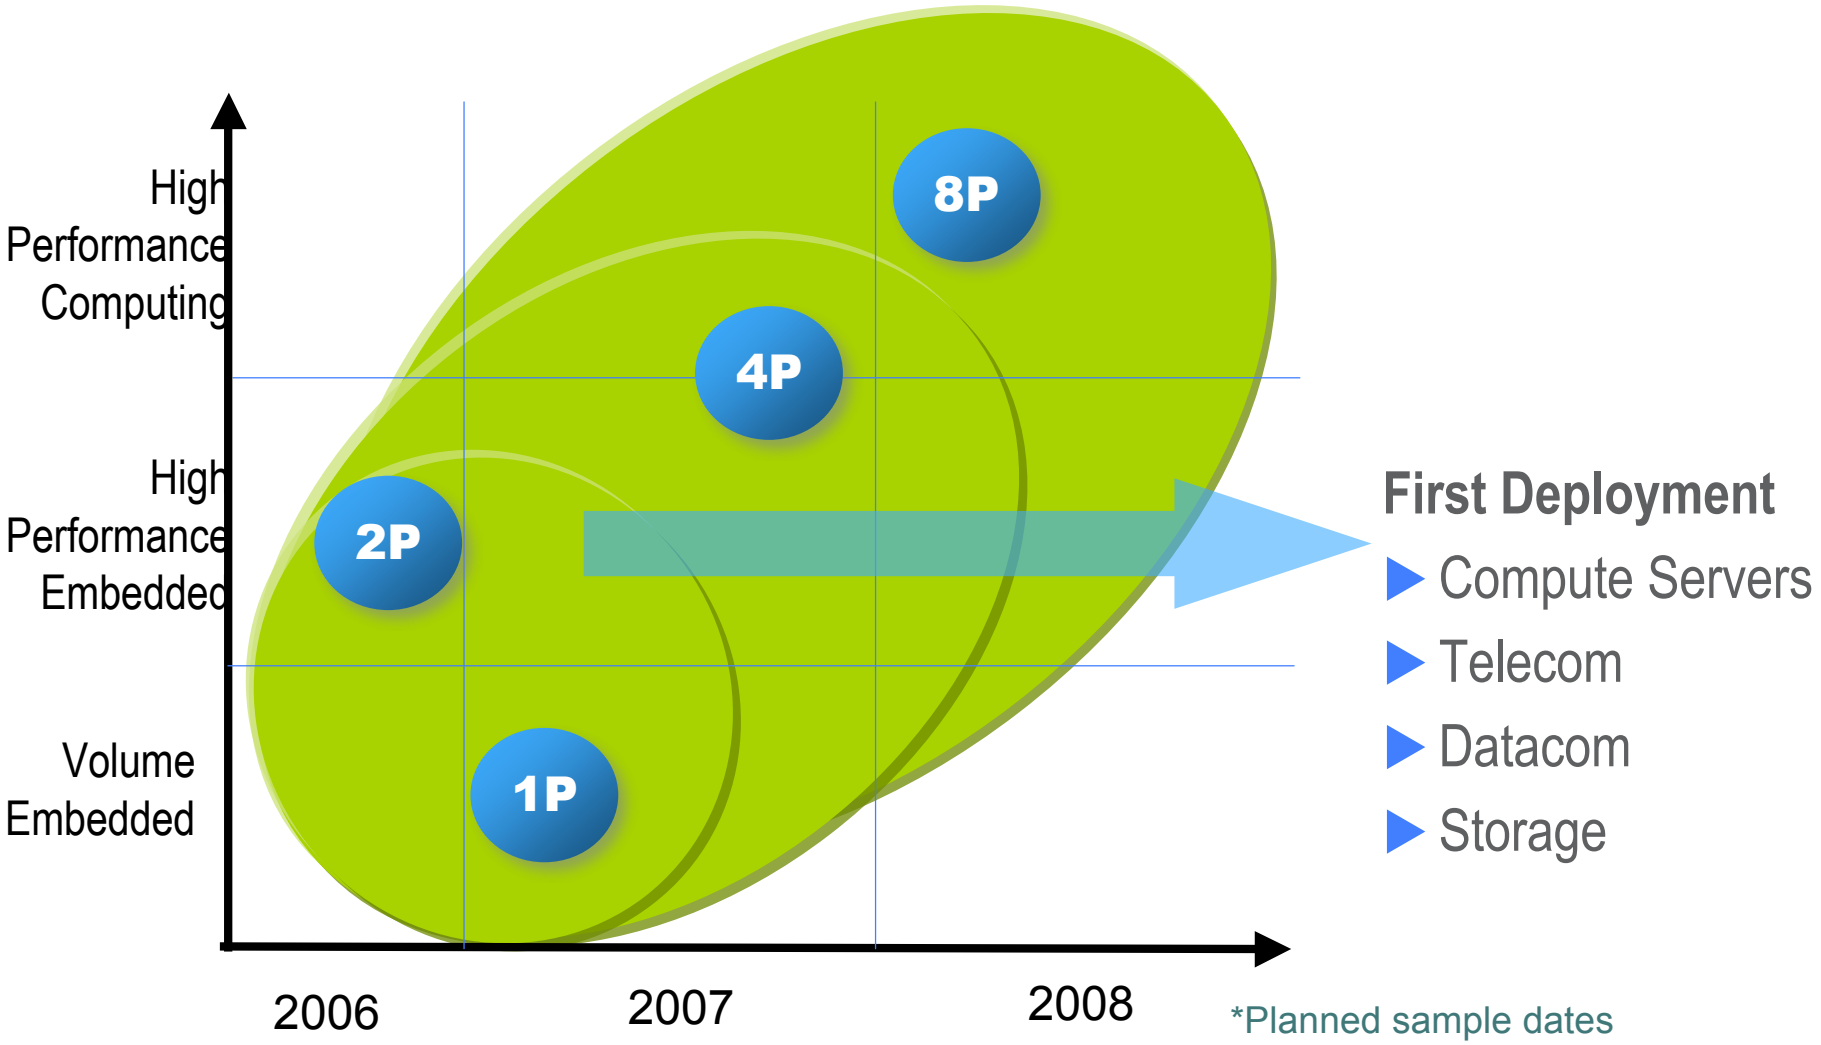

The Core

PA SEMI

Power to Perform™

PA6T Core

- **Fully compliant Power implementation**
  - Power 2.04 architecture spec
  - Full FPU and VMX SIMD support
  - 32- and 64-bit
  - Big- and little-endian modes
    - Bi-endian per process
- **Super-scalar, out-of-order design**
  - Quad fetch, triple issue
  - 64-entry scheduler
  - Strongly ordered memory model
- **Hypervisor and virtualization support**
- **High performance memory hierarchy @ 2Ghz**
  - L1 data—32GB read or write
  - L2 data—16GB read plus 16GB write
  - Memory DDR2-1066 —16GB read or write
  - 16 transactions in flight
- **CONEXIUM Interchange**
  - 1G transactions per second
  - 64GB/sec data peak
- **Extensive power management capabilities**

Processor Pipeline
Multi-Issue, Out-of-Order, Superpipelined

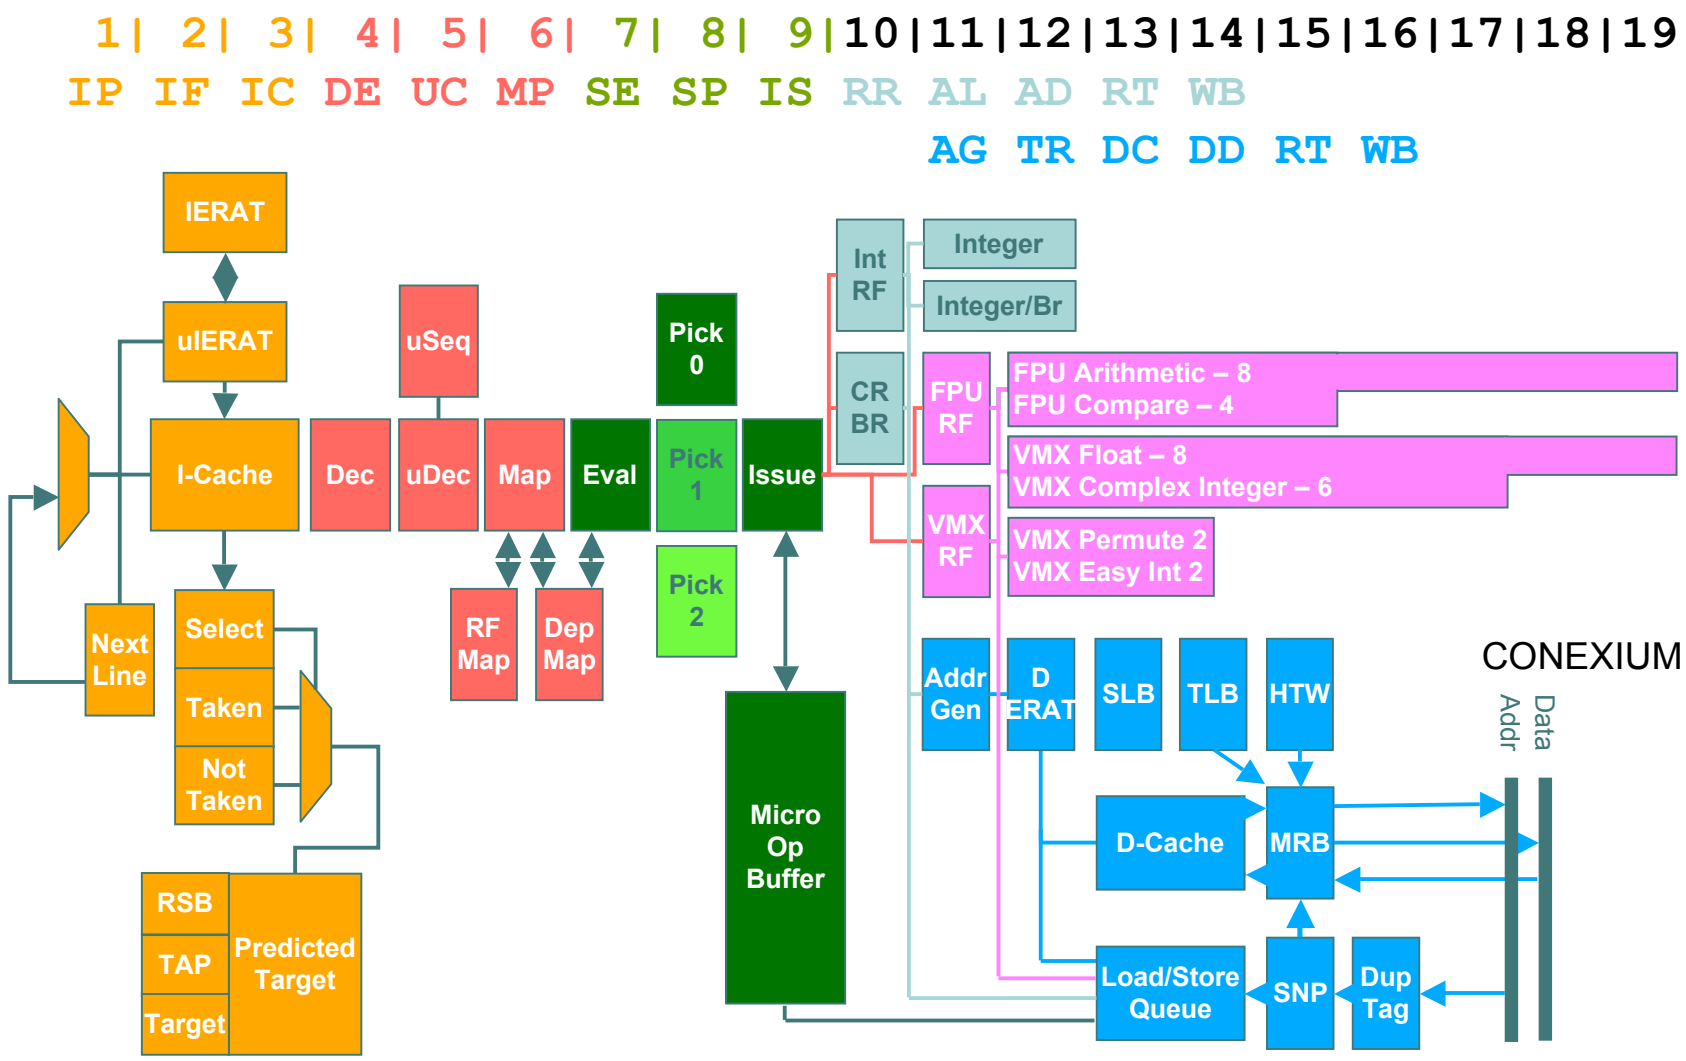

Processor Pipeline — Front End

1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
IP IF IC DE UC MP SE SP IS

128-entry IERAT
4-entry uIERAT

Icache 64K bytes
2-way associative
4 instruction/fetch

Predictors
16-entry next line
32K-bit conditional branch history

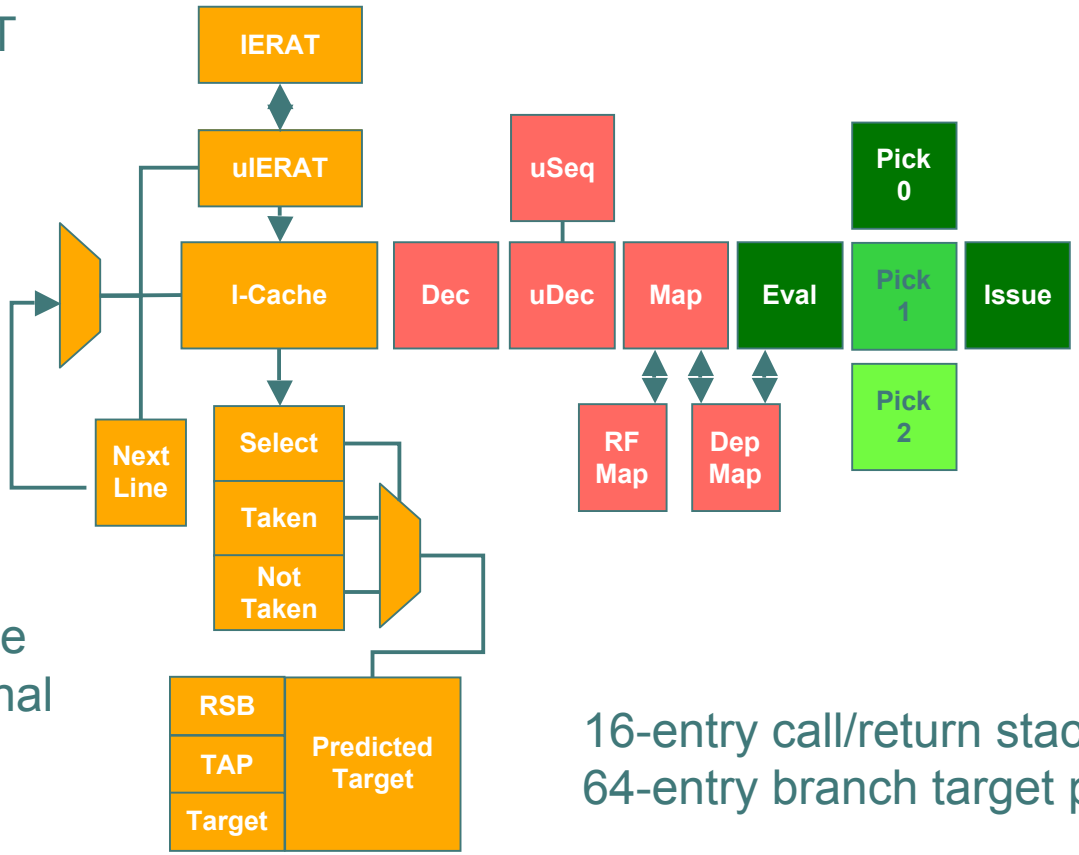

16-entry call/return stack
64-entry branch target predictor

Processor Pipeline — Load/Store Processing

Latency - Load to use

L1 cache	4
L2 cache	22
Open page	90
Closed page	110
Remote L1	30

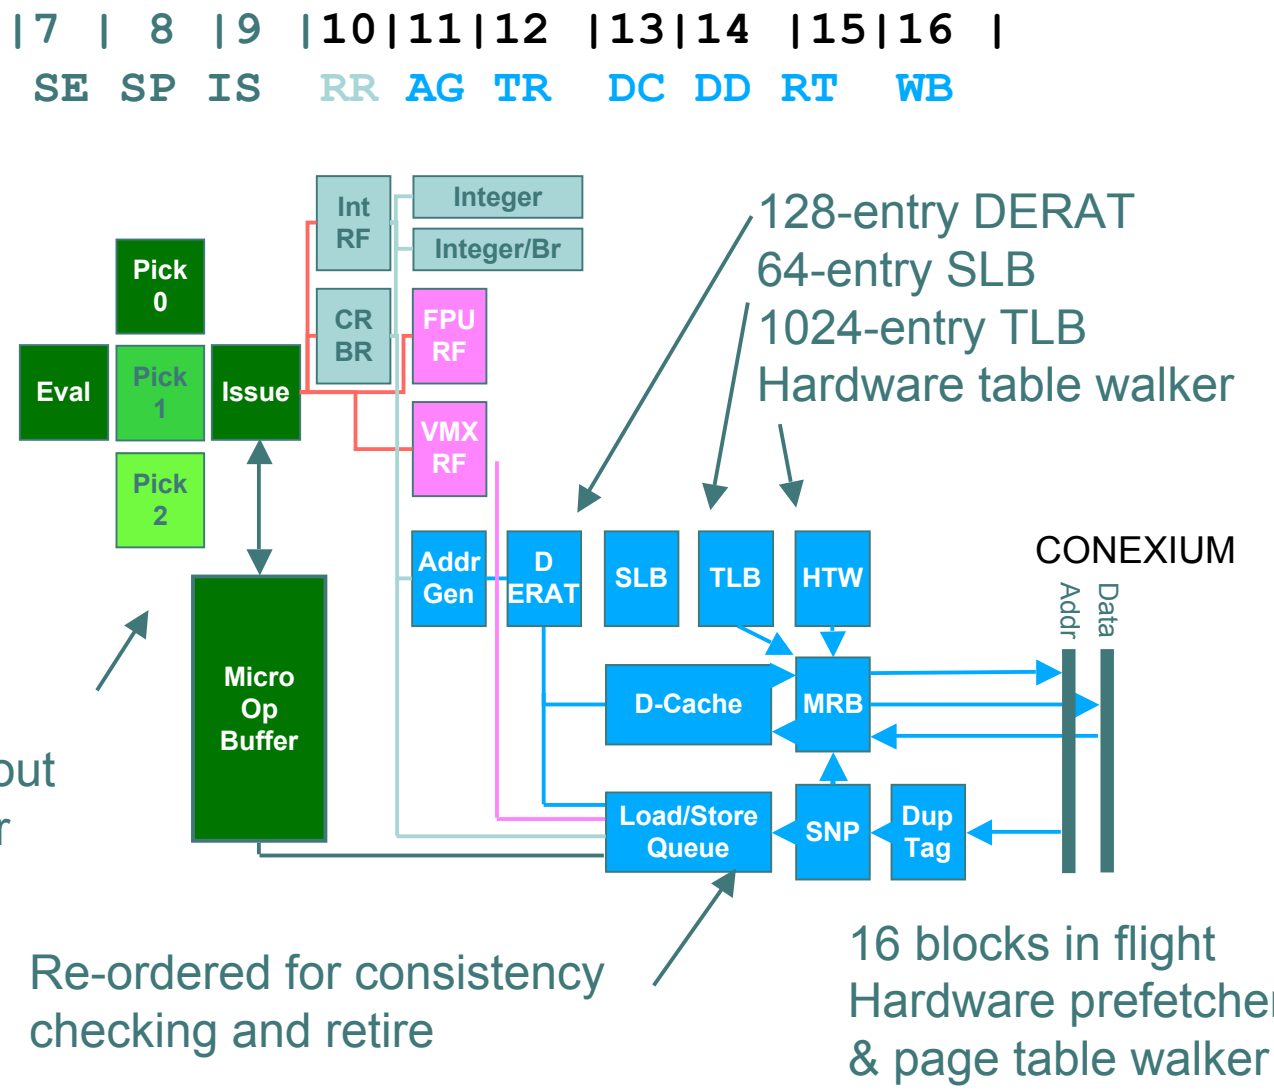

PWRficient Platform Architecture

PA SEMI

Power to Perform™

PWRficient's Interconnects

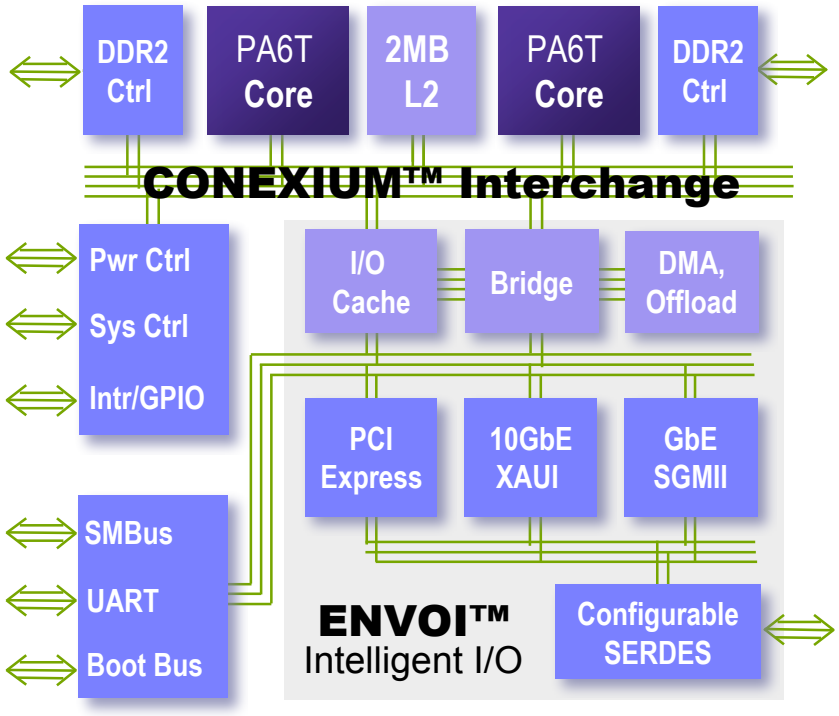

- CONEXIUM connects coherent agents
  - Core(s), L2 cache(s), memory controllers, and I/O bridge
- ENVOI connects I/O
  - Bridge handles addressed traffic
  - DMA handles packet traffic
  - Protocol engines can be either type
    - PCI Express—addressed traffic
    - XAUI & SGMII MACs—packet traffic
- Configurable SERDES connects with any protocol engine (Ethernet MACs, PCI Express)

CONEXIUM Interconnect Details

- **CONEXIUM Interchange: Coherent, Ordered, X-Bar**
  - Transaction initiators: core, I/O bridge
  - All devices respond
  - MOESI-style protocol

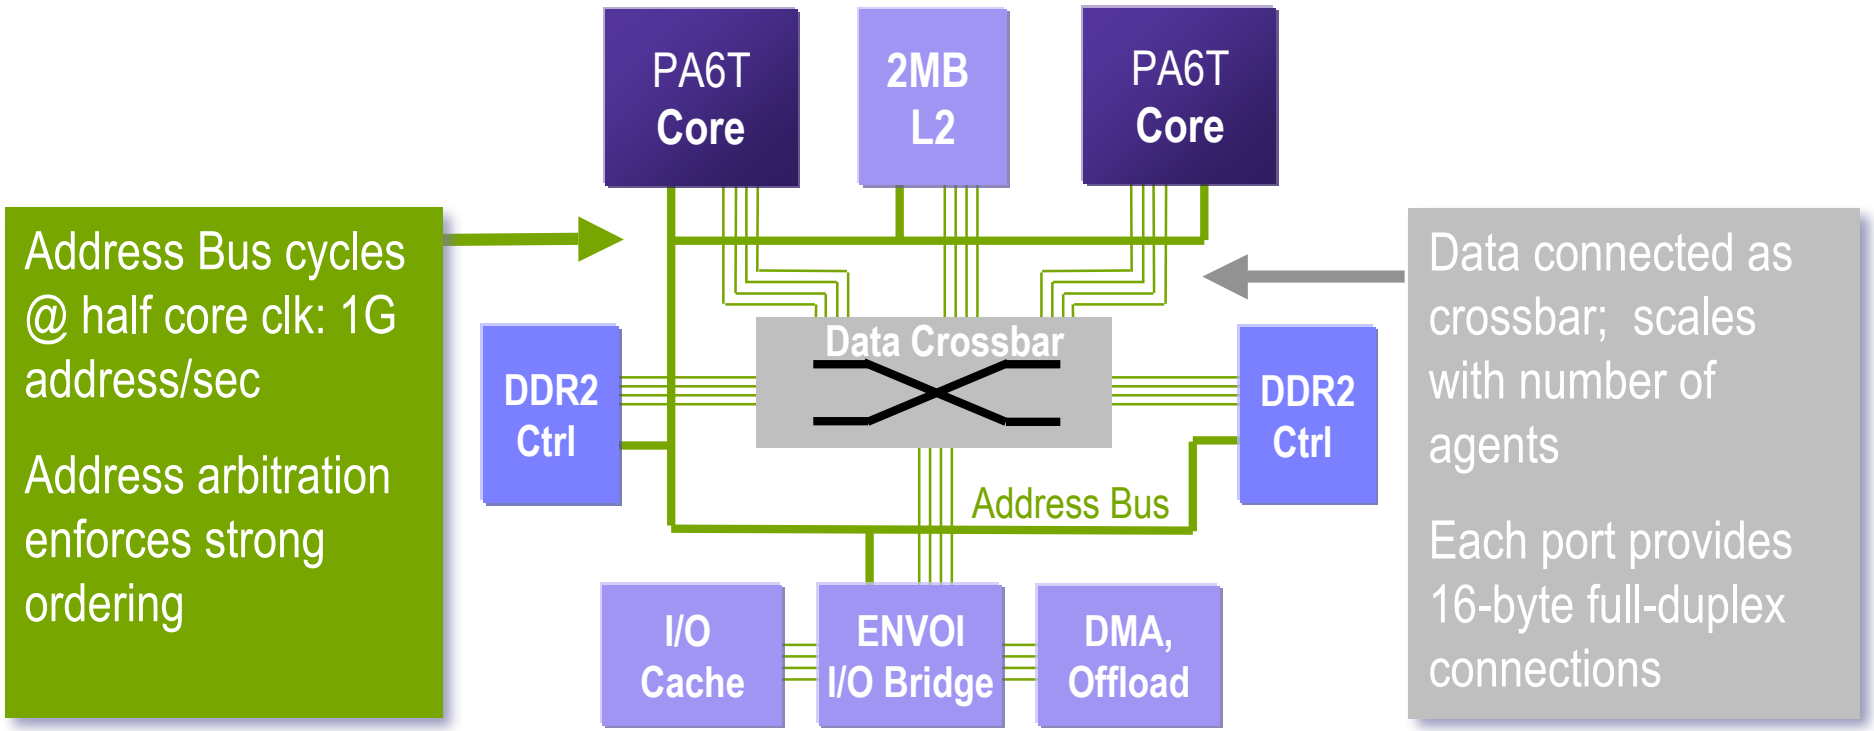

ENVOI — Intelligent I/O

- **Interoperability between PCI Express, packets, DMA, and memory**
  - Centralized DMA model
  - Global arbitration for bandwidth allocation and fairness
  - Address Translation

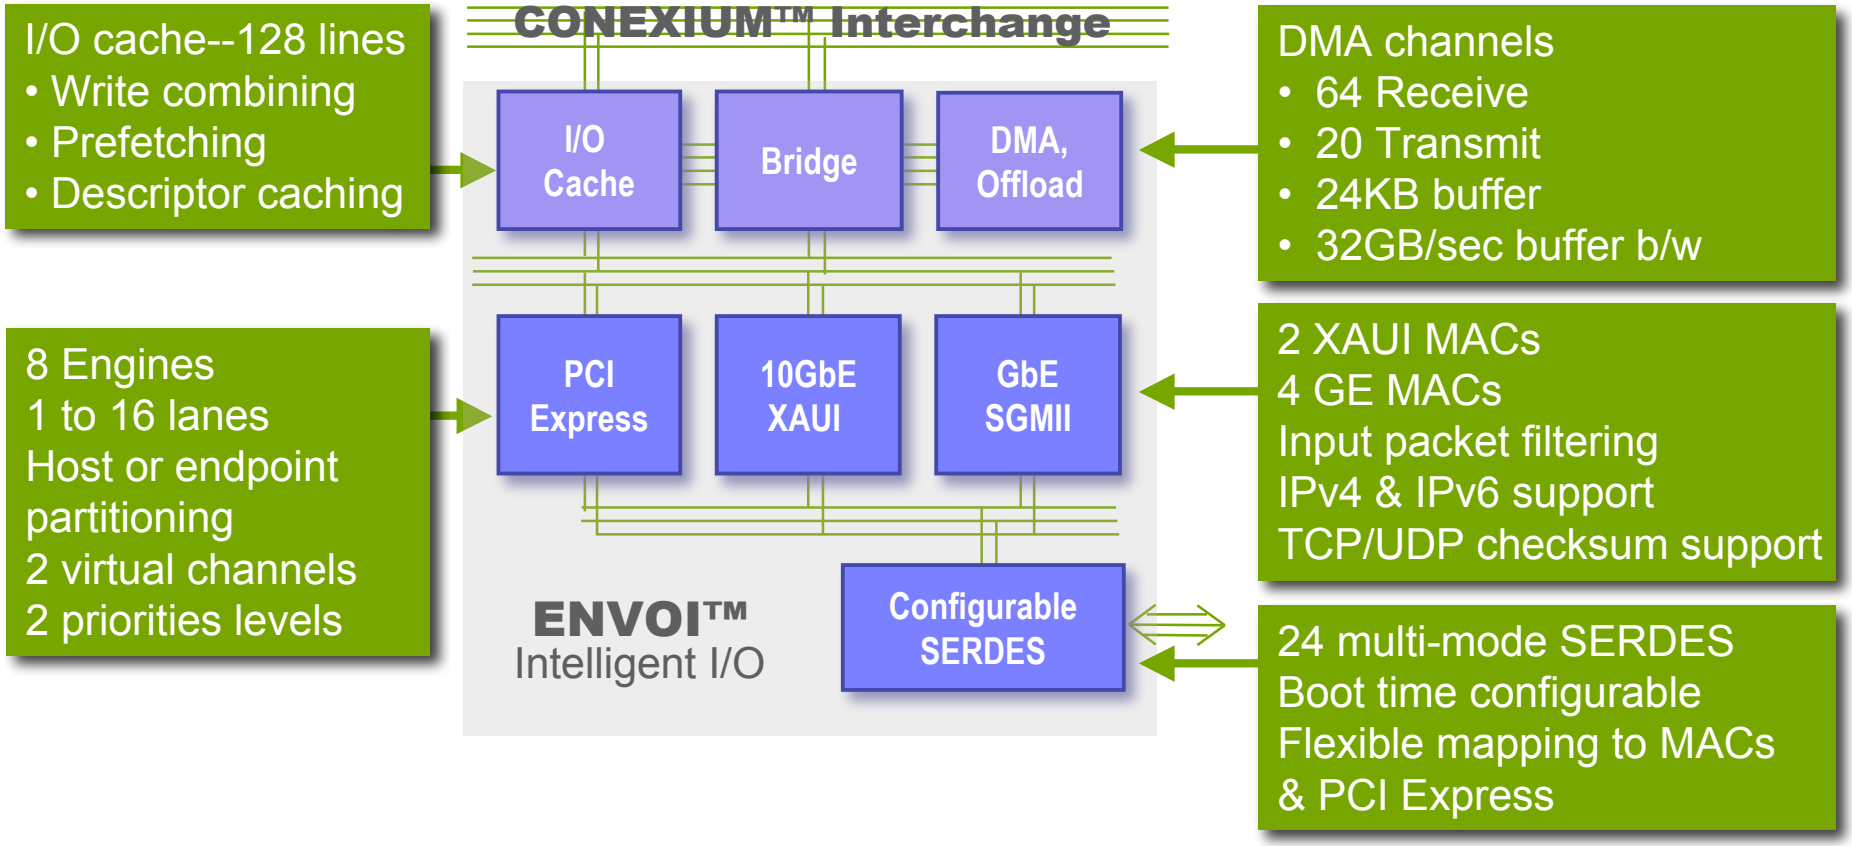

ENVOI DMA and Offload Engines

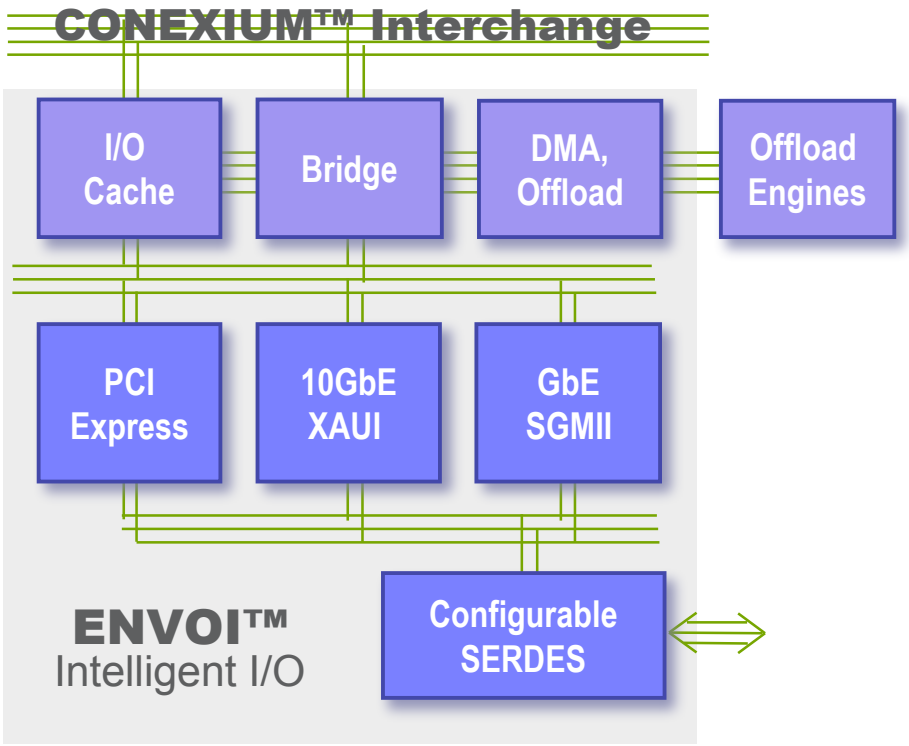

- Memory, PCI Express, MACs, or offload engines can be source or target
  - Scatter-gather
  - Memory-to-memory
  - I/O-to-I/O, I/O-to-Memory, Memory-to-I/O
- Offload engines transform data
  - Encryption
    - Bulk: 3DES, AES, ARC4, Kasumi (f8)
    - Signature: MD5, SHA-1, SHA-256, Kasumi (f9)
    - Packet: IPSEC, SSL
  - CRC
    - General-purpose engine
  - Checksums
    - TCP, etc
  - RAID XOR engine

ENVOI Platform I/O

- **Power & System Controller**
  - Power and frequency management
  - Baseband management controller interface
  - Watchdog and regular timers
  - Reset control
- **Interrupt Controller**
  - OpenPIC
  - Multiprocessor interrupt dispatch
  - Supports 512 PCI Express MSIs
  - Low-latency interrupts & status
- **Boot / local bus**
  - LPC protocol -direct connect for boot flash
    - I/O space and DMA modes
  - Multiplexed AD protocols—IDE support
    - NAND flash and CompactFlash support
- **Platform Serial I/O**
  - Multiple SMBus channels
    - Master and slave
    - Software-defined protocols
  - Multiple UARTs

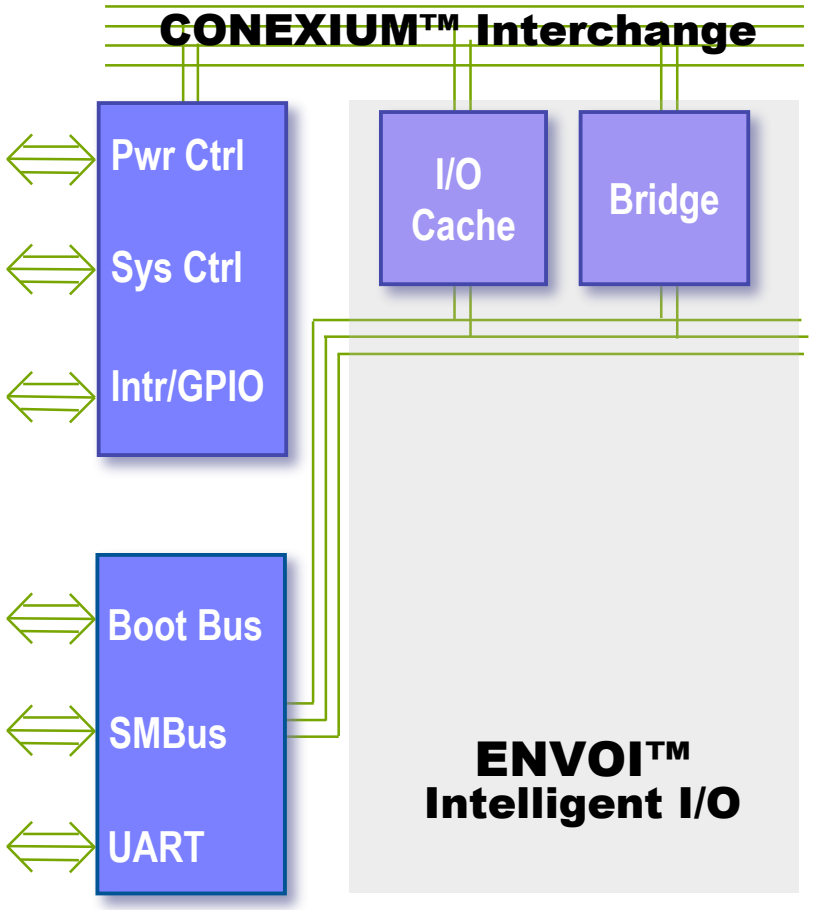

PWRficient Performance

PA SEMI

Power to Perform™

Power Efficient

Power is first-order design principle

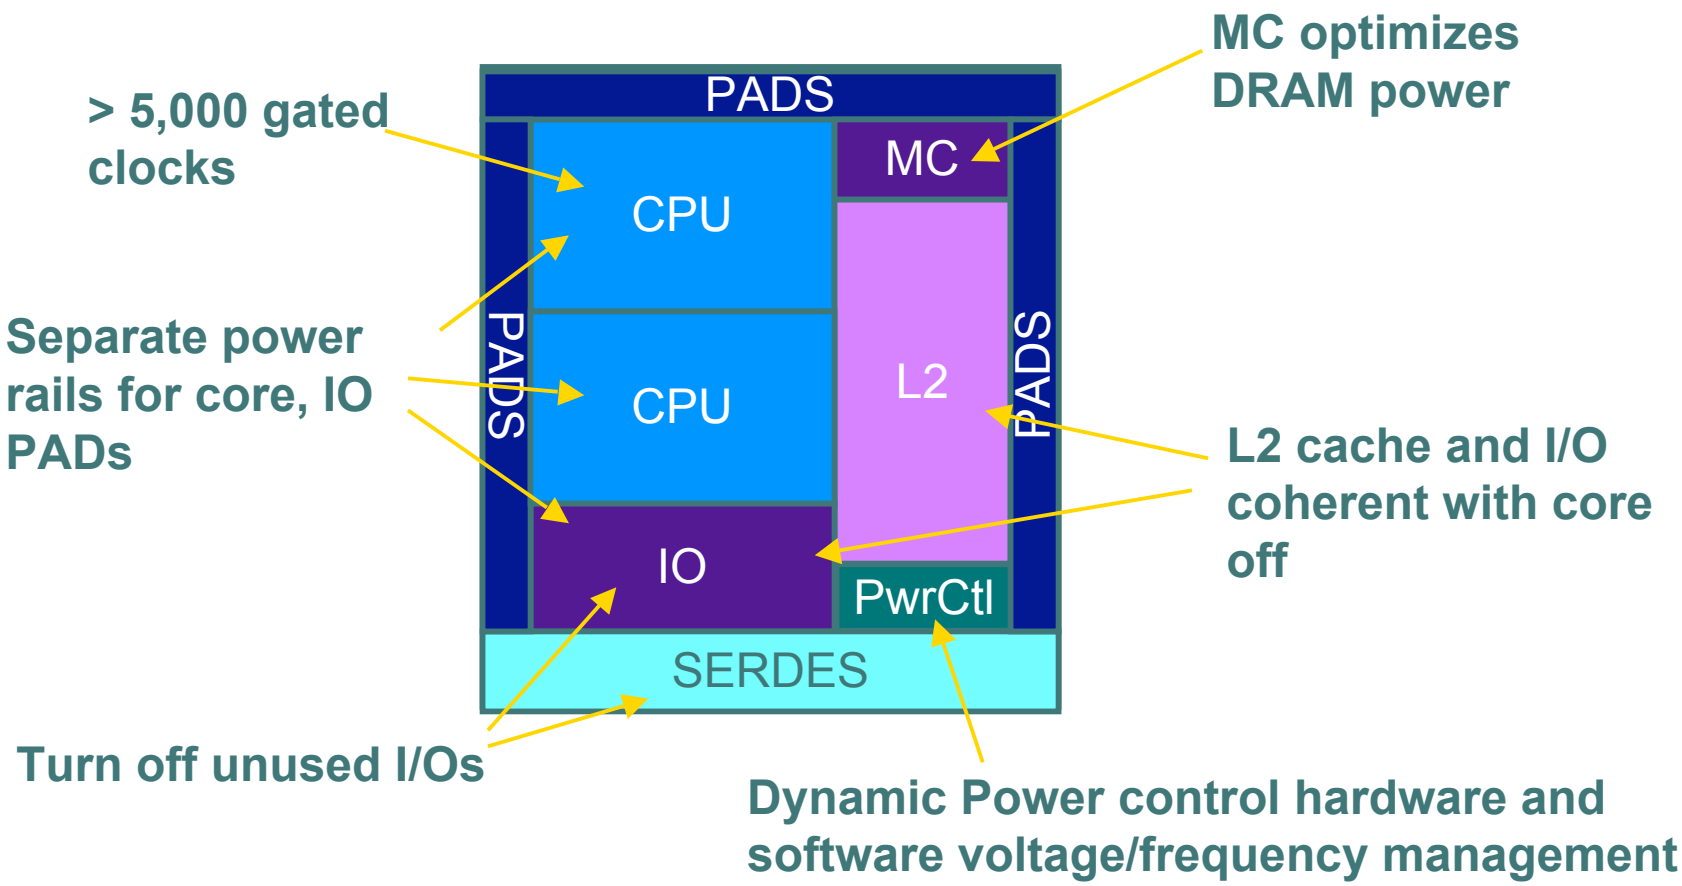

Power Summary

	Typical Power	Worst Case Power	Comments
PA6T Core @ 2Ghz	4W	7W	Core-only power
Platform, dual-core @ 2GHz	13W	25W	Full SOC, everything on
Platform, dual-core @ 1.5Ghz	8W	15W	Full SOC @ 1.5GHz core
Platform, dual-core @ 1Ghz	6W	10W	Full SOC @ 1GHz core
Platform, nap mode	1W	2W	Coherent L2 & IO

High-Performance At Low-Power Across A Range Of Metrics

PWRficient 1682M provides mainstream performance at low power

General-purpose computing

- SPECint™ >1000 per core*

Floating-point performance

- SPECfp™ >2000 per core*
- Imaging
  - FFT 24 GFlops/sec (total)*

System bandwidth

- Memory-to-memory copy
  - 10 Gigabytes/sec*
- High-speed serial I/O
  - 24 SERDES for total I/O bandwidth up to 104Gbps peak

Application offloads

- TCP/IP termination
  - > 20Gbps†
- Encryption
  - 10Gbps bulk IPSec and SSL encryption†
  - 3,000 public-key handshakes/sec in software*
- Storage
  - 2.0GB/s RAID5 (data+parity)†

*Estimated max sustained performance at 2GHz
†Estimated peak performance at 2GHz

Example System Block Diagrams

PA SEMI

Power to Perform™

Exception Processor

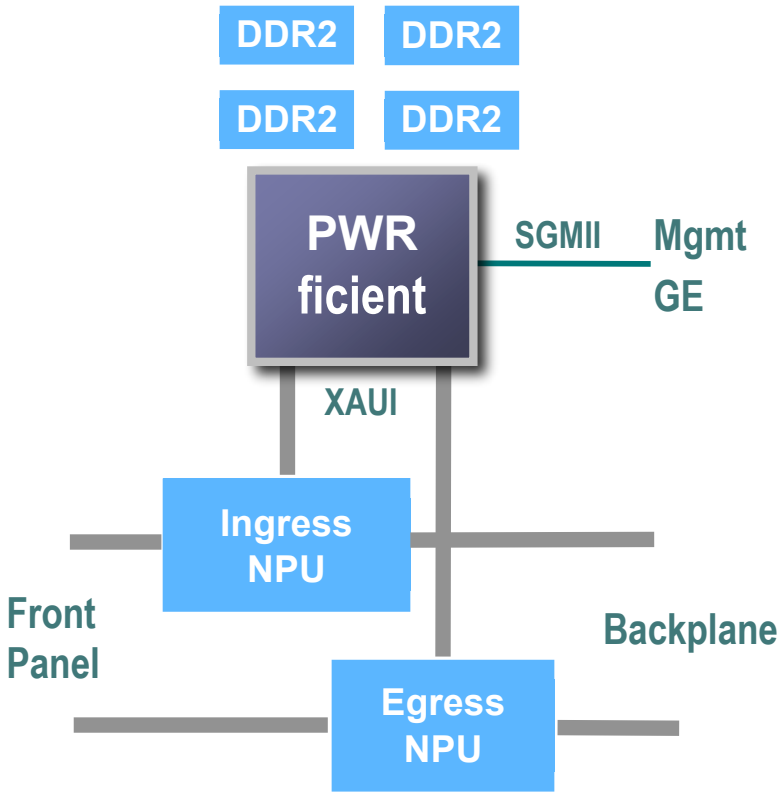

- ASICs/NPUs cover simple packets
- Interesting packets to PWRficient
- XAUI or PCI Express interconnect
  - Versatile “connectivity”
  - Bursts at 10Gbps
    =>No buffer needed in NPU
- PWRficient processing
  - Scalable as feature set changes, new protocols, attacks...
  - High performance => long product life

Embedded Processing Platform

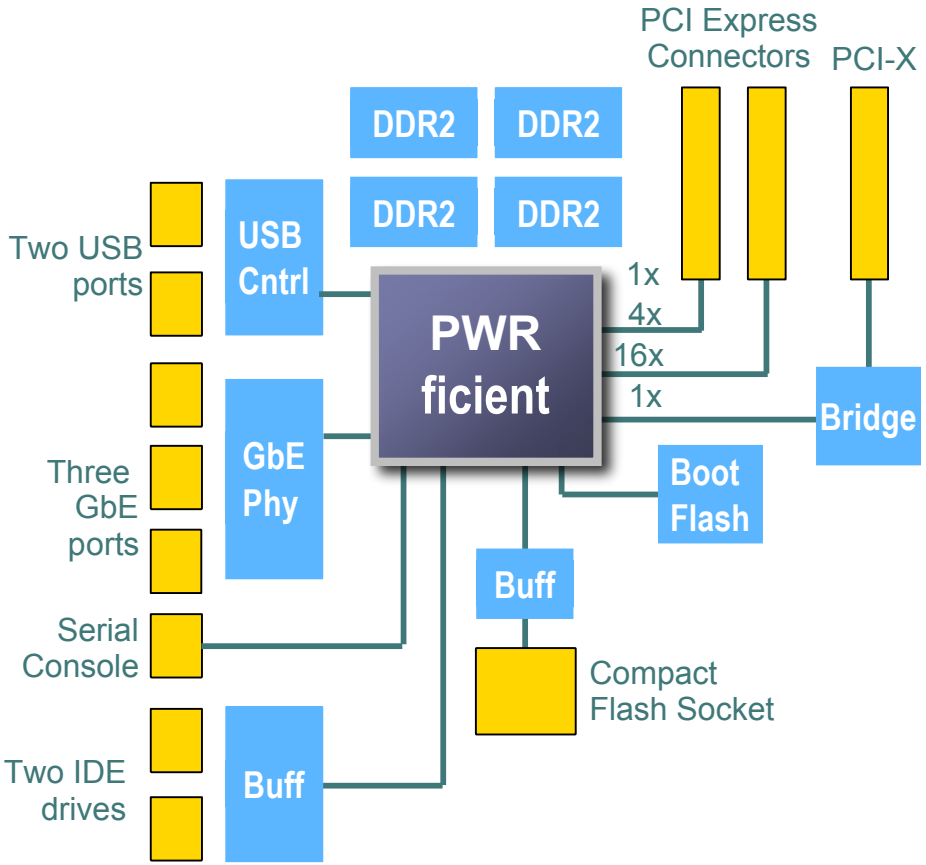

- **Versatile embedded platform**
  - Standard I/O connectors
    - PCI Express (4-lane, 16-lane)
    - PCI-X
    - IDE
    - USB
  - Three GbE network interfaces
- **Self-contained Linux appliance**
  - Flow-through GbE applications
    - For example, security appliance
    - Internal, external, and DMZ networks
- **High-performance, low power**
  - Full range of power management modes
  - Enables compact, fanless form factors

High-Density HPC/Server

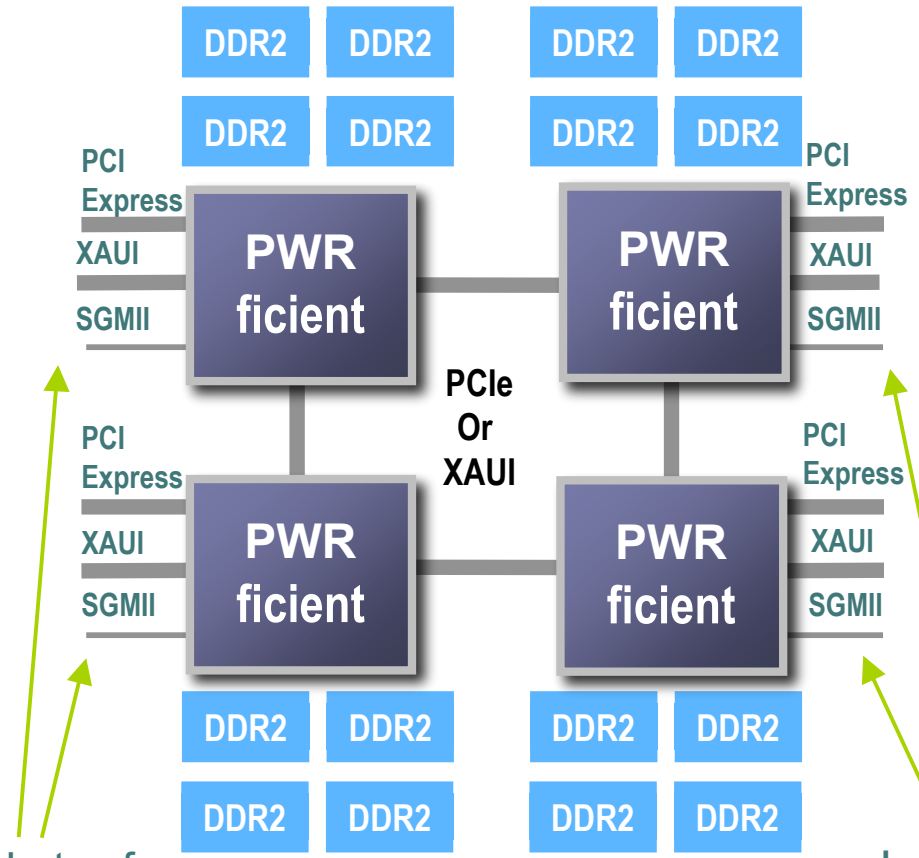

- **The 100W solution**
  - 8 Cores ~50 W
  - ~50 Watt for DRAM & I/O
  - Low power = high compute density
- **High Performance**
  - 32Gflops/sec double precision
  - 128Gflops/sec single precision
- **High end Memory & IO system**
  - Large memory support – 128GBytes
  - 16 ranks per SOC
  - 64 GB/sec memory bandwidth
  - 96 SERDES @ 2.5GHz

Summary

- **New Technology and I/O enable new possibilities**
  - Optimize performance, power, latency, density
- **From-scratch design maximizes the gains**
- **Modular design gives flexible chip options**
- **Interesting convergence of needs across a wide range of design points**
  - Networking, telecom, servers, imaging
- **Performance and power enable wide range of applications**

Contact P. A. Semi

- For further information, please visit the P. A. Semi website at www.pasemi.com
- Kindly direct sales inquiries to sales@pasemi.com
- Full contact information:

P.A. Semi, Inc.

3965 Freedom Circle, Floor 8

Santa Clara

CA 95054-1203 USA

Main: 408.200.4500

Fax: 408.200.4501

Thank You

PA SEMI

Power to Perform™